Modeling and Estimation of Discrete-Time Reciprocal Processes via Probabilistic Graphical Models

Francesca Paola Carli

arXiv:1603.04419v3 [stat.ML] 13 May 2016

Abstract

Reciprocal processes are acausal generalizations of Markov processes introduced by Bernstein in 1932. In the literature, a significant amount of attention has been focused on developing *dynamical models* for reciprocal processes. In this paper, we provide a *probabilistic graphical model* for reciprocal processes. This leads to a principled solution of the smoothing problem via *message passing algorithms*. For the finite state space case, convergence analysis is revisited via the *Hilbert metric*.

I. Introduction

Non causal random processes arise in many areas of science and engineering. For these processes, the index set usually represents space instead of time. The class of non causal reciprocal processes was introduced by Bernstein in 1932 [3] and studied by many authors [20], [21], [22], [10], [26], [28], [25], [8], [9], [12], [40]. A $\mathbb{R}^n$–valued stochastic process $\mathbf{X}_k$ defined over the interval $\mathcal{I} = [0, N]$ is said to be reciprocal if for any subinterval $[K, L] \subset \mathcal{I}$, the process in the interior of $[K, L]$ is conditionally independent of the process in $\mathcal{I} - [K, L]$ given $\mathbf{X}_K$ and $\mathbf{X}_L$. Reciprocal processes are a natural generalization of Markov processes: from the definition it immediately follows that Markov processes are necessarily reciprocal, but the converse is not true [20]. Moreover multidimensional Markov random fields reduce in one dimension to a reciprocal process, not to a Markov process. To attest the relevance of reciprocal processes from an engineering point of view, note, for example, that the steady-state distribution

Francesca Paola Carli is with the Department of Engineering, University of Cambridge, United Kingdom, fpc23@cam.ac.uk

of the temperature along a heated ring or a beam subjected to random loads along its length can be modeled in terms of a reciprocal process. Applications to tracking of a ship-trajectory [13], estimation of arm movements [36], and synthesis of textured images [34] have also been considered in the literature.

Starting with Krener's work [26], a significant amount of attention has been focused on developing *dynamical models* for reciprocal processes. Both the continuous and discrete–time case have been addressed. In this paper, our focus is on discrete–time reciprocal processes. In [28] it has been shown that a discrete–time *Gaussian* reciprocal process admits a second–order nearest–neighbor model driven by a locally correlated noise, where the noise correlation structure is specified by the model dynamics. This model recalls state–space models for Markov processes but is *acausal* (the system does not evolve recursively in the direction of increasing or decreasing values of k) and the driving noise is not white. Second order state space models for discrete–time *finite–state* reciprocal processes have been derived in [12] (see also [11]).

In this paper, we provide a *probabilistic graphical model* for reciprocal processes with cyclic boundary conditions. In particular, it is shown that a reciprocal process with cyclic boundary conditions admits a single loop undirected graph as a perfect map. This approach is *distribution–independent* and leads to a principled solution of the *smoothing problem* via belief propagation (a.k.a. sum–product) algorithms. In this scheme, the estimated posteriors ("beliefs") are computed as the product of incoming messages at the corresponding node, messages being updated through local computations (every given node updates the outgoing messages on the basis of incoming messages at the previous iteration alone). For tree-structured graphs, the sum–product algorithm is guaranteed to converge to the correct posterior marginal [32]. Nevertheless, since message passing rules are purely local, the sum–product algorithm can also be applied to loopy networks as an approximation scheme. As mentioned above, the graphical model associated to a reciprocal process is a single–loop network, which is not a tree. Convergence of sum–product algorithms for single–loop networks has been studied in the literature (see [38], [39] and references therein). For the finite state space case, we revisit convergence analysis via the Hilbert metric. This approach is geometric in nature, leveraging on contraction properties of positive operators that map a quite generic cone into itself, and as such it can be extended to analyze convergence of message passing algorithms in more general settings (state–spaces, see the companion paper [7], and graph topologies), thus providing a unifying framework for the analysis of convergence of message

passing algorithms for a single loop undirected graph, that has instead been treated via ad hoc arguments in the literature (see e.g. [38], [39] where different techniques has been employed for the Gaussian and the finite state space cases). To recap, the contribution of the paper is threefold: (i) providing a probabilistic graphical model for reciprocal processes; (ii) solving the smoothing problem via message passing algorithms; (iii) providing an alternative analysis of convergence of such algorithms leveraging on contraction properties of positive operators with respect to the Hilbert metric.

The paper is organized as follows. In Section II reciprocal processes are introduced. Second–order nearest neighbor models for reciprocal processes are discussed in Section III. Section IV reviews relevant theory about probabilistic graphical models. The probabilistic graphical model associated to a reciprocal process with cyclic boundary conditions is derived in Section V where it is shown that a reciprocal process with cyclic boundary conditions admits a single–loop Markov network as perfect-map. The smoothing problem for reciprocal processes is solved in Section VI via loopy belief propagation. Sections VII and VIII introduce the Hilbert metric and discuss its relevance for stability analysis of linear positive systems. Contraction properties of positive operators with respect to the Hilbert metric are exploited to prove convergence of loopy belief propagation for finite state reciprocal processes in Section IX. Section X ends the paper.

II. Reciprocal Processes

A stochastic process $\mathbf{X}_t$ defined on a time interval $\mathcal{I}$ is said to be *Markov* if, for any $t_0 \in \mathcal{I}$, the past and the future (with respect to t_0) are conditionally independent given $\mathbf{X}_{t_0}$. A process is said to be *reciprocal* if, for each interval $[t_0, t_1] \subset \mathcal{I}$, the process in the interior of $[t_0, t_1]$ and the process in $\mathcal{I} - [t_0, t_1]$ are conditionally independent given $\mathbf{X}_{t_0}$ and $\mathbf{X}_{t_1}$. Formally [21]

Definition 2.1: A $(\mathcal{X}, \Sigma)$–valued stochastic process $\{\mathbf{X}_t\}$ on the interval $\mathcal{I}$ with underlying probability space $(\Omega, \mathcal{A}, P)$ is said to be *reciprocal* if

$$P(AB \mid \mathbf{X}_{t_0}, \mathbf{X}_{t_1}) = P(A \mid \mathbf{X}_{t_0}, \mathbf{X}_{t_1}) P(B \mid \mathbf{X}_{t_0}, \mathbf{X}_{t_1}), \tag{1}$$

$\forall t_0 < t_1$, $[t_0, t_1] \subset \mathcal{I}$, where A is the σ–field generated by the random variables $\{\mathbf{X}_r : r \notin [t_0, t_1]\}$ and B is the σ-field generated by $\{\mathbf{X}_r : r \in (t_0, t_1)\}$.

From the definition we have that Markov processes are necessarily reciprocal, while the converse is generally not true [20]. The class of reciprocal processes is thus larger than the

Markov class, and it naturally extends to the multidimensional case where the parameter set of the process is not linearly ordered. In fact multidimensional Markov random fields reduce in one dimension to a reciprocal process, not to a Markov process.

In this paper, we consider reciprocal processes defined on the discrete circle $\mathbb{T}$ with $N+1$ elements $\{0, 1, \ldots, N\}$ (which corresponds to imposing the cyclic boundary conditions $\mathbf{X}_{-1} = \mathbf{X}_N$, $\mathbf{X}_{N+1} = \mathbf{X}_0$, see [28], [35] and Section III below) so that the additional conditional independence relations

$$\mathbf{X}_0 \perp\!\!\!\perp \{\mathbf{X}_2, \ldots, \mathbf{X}_{N-1}\} \mid \{\mathbf{X}_1, \mathbf{X}_N\} \,,$$

$$\mathbf{X}_N \perp\!\!\!\perp \{\mathbf{X}_1, \ldots, \mathbf{X}_{N-2}\} \mid \{\mathbf{X}_0, \mathbf{X}_{N-1}\}$$

hold.

Starting with Krener's work [26], a significant amount of attention has been focused on developing *dynamical models* for reciprocal processes. In this paper, our focus is on discrete–time reciprocal processes. In the next Section we briefly review dynamical models for discrete–time reciprocal processes, that were first introduced in [28]. In Section V we provide a *probabilistic graphical model* representation of reciprocal processes.

III. Second–order Models of Reciprocal Processes

Let $\mathbf{X}_k$ be a zero-mean process defined over the finite interval $\mathcal{I} = [0, N]$ and taking values in $\mathbb{R}^n$. It is well–known that if $\mathbf{X}_k$ satisfies the recursion equation

$$\mathbf{X}_{k+1} = \mathbf{A}_k \mathbf{X}_k + \mathbf{W}_k \tag{2}$$

where $\mathbf{W}_k$ is a zero–mean random process with

$$\mathbb{E}\left[\mathbf{W}_k \mathbf{W}_l^\top\right] = \mathbf{I}\,\delta_{kl} \tag{3}$$

and $\mathbf{X}_0$ is a zero–mean random variable such that

$$\mathbb{E}\left[\mathbf{W}_k \mathbf{X}_0^\top\right] = 0 \tag{4}$$

then $\mathbf{X}_k$ is Markov. If $\mathbf{X}_k$ is Gaussian, then the converse is also true, namely it can be shown (see e.g. [1]) that a Gaussian process is Markov if and only if it satisfies (2) with noise structure (3), (4).

For a reciprocal process, the following holds. Let $1 \leq k \leq N-1$, and consider the model

$$-\mathbf{M}_k^- \mathbf{X}_{k-1} + \mathbf{M}_k^0 \mathbf{X}_k - \mathbf{M}_k^+ \mathbf{X}_{k+1} = \mathbf{E}_k \,, \tag{5}$$

where $\mathbf{M}_k^0$, $\mathbf{M}_k^+$, $\mathbf{M}_k^-$ are such that

$$\mathbf{M}_k^0 = (\mathbf{M}_k^0)^\top, \qquad \mathbf{M}_k^+ = (\mathbf{M}_{k+1}^-)^\top \tag{6}$$

and the driving noise $\mathbf{E}_k$ satisfies

$$\mathbb{E}\,[\mathbf{E}_k \mathbf{X}_l^\top] = \mathbf{I}\,\delta_{kl} \tag{7}$$

and is locally correlated with covariance $\mathbf{\Sigma}_e$

$$[\mathbf{\Sigma}_e]_{k,l} = \begin{cases} \mathbf{M}_k^0, & \text{for } l = k \\ -\mathbf{M}_k^+ & \text{for } l = k+1 \\ 0 & \text{otherwise}\,. \end{cases} \tag{8}$$

Equations (5)–(8) specify a second–order nearest–neighbor model. The model recalls standard first–order state–space models for Markov processes but it is *acausal* (the system does not evolve recursively in the direction of increasing or decreasing values of k). Also, the driving noise $\mathbf{E}_k$ is not white, but locally correlated. Notice that, in order to completely specify $\mathbf{X}_k$ over the interval $\mathcal{I} = [0, N]$, some boundary conditions must be provided. Following [28], in this paper we consider *cyclic boundary conditions*, namely we assume

$$\mathbf{X}_{-1} = \mathbf{X}_N, \qquad \mathbf{X}_{N+1} = \mathbf{X}_0\,. \tag{9}$$

These conditions are equivalent to extending cyclically the model (5) and the noise structure (7), (8) to the whole interval $\mathcal{I} = [0, N]$, provided that, in these identities, $k-1$ and $k+1$ are defined modulo $N+1$. Equation (5) with cyclic boundary conditions (9) can be written in matrix form as

$$\mathbf{M}\mathbf{X} = \mathbf{E} \tag{10}$$

where

$$\mathbf{X} = \begin{bmatrix} \mathbf{X}_0 \\ \mathbf{X}_1 \\ \vdots \\ \mathbf{X}_N \end{bmatrix}, \quad \mathbf{E} = \begin{bmatrix} \mathbf{E}_0 \\ \mathbf{E}_1 \\ \vdots \\ \mathbf{E}_N \end{bmatrix},$$

and matrix $\mathbf{M}$ given by

$$\mathbf{M} = \begin{bmatrix} \mathbf{M}_0^0 & -\mathbf{M}_0^+ & 0 & \dots & 0 & -\mathbf{M}_0^- \\ -\mathbf{M}_1^- & \mathbf{M}_1^0 & -\mathbf{M}_1^+ & 0 & \dots & 0 \\ \dots & & & & & \dots \\ 0 & \dots & 0 & -\mathbf{M}_{N-1}^- & \mathbf{M}_{N-1}^0 & -\mathbf{M}_{N-1}^+ \\ -\mathbf{M}_N^+ & 0 & \dots & 0 & -\mathbf{M}_N^- & \mathbf{M}_N^0 \end{bmatrix}. \tag{11}$$

It can be shown that a process $\{\mathbf{X}_k\}$ satisfying (5)–(9) is reciprocal. Moreover, if the process is Gaussian, the converse is also true. To be more precise:

Theorem 3.1: [28] Let $\mathbf{X}_k$ be a zero–mean *Gaussian* process on $\mathbb{T}$ whose covariance $\mathbf{\Sigma}_x$ is nonsingular, i.e. $\mathbf{\Sigma}_x > 0$. Then $\mathbf{X}_k$ is *reciprocal* if and only if it admits a well–posed second–order descriptor model of the form (5)–(9).

State space modeling for *finite state space* reciprocal processes has been separately addressed in [12] (see also [11]). While different state space models have been proposed in the literature for the Gaussian [28] and finite state space [12] cases, the probabilistic graphical model we introduce in Section V is distribution independent. The following subsection provides a characterization of Gaussian reciprocal processes in terms of the sparsity pattern of their precision matrix.

Characterization via Covariance Matrix

If the $\mathbf{X}_k$'s are *normally distributed*, an important characterization in terms of sparsity pattern of the inverse of the covariance matrix (a.k.a. the precision matrix) holds. To start, let's recall the following (see, e.g., [15], [27]):

Theorem 3.2: The (i, j)–th (block)–entry of the inverse covariance matrix is zero if and only if the i–th and j–th (vector)–components of the underlying Gaussian random vector are conditionally independent given the other (vector)–components.

Now, it is well known that $\mathbf{\Sigma}_x > 0$ is the covariance of a (vector–valued) Markov process if and only if $\mathbf{\Sigma}_x^{-1}$ is (block) tridiagonal (see [1]). In [28] the following characterization of nonsingular Gaussian reciprocal processes on a finite interval was obtained.

Theorem 3.3: $\mathbf{\Sigma}_x > 0$ is the covariance matrix of the Gaussian reciprocal process (10) if and

only if $\mathbf{\Sigma}_x^{-1}$ has the block tridiagonal structure

$$\mathbf{\Sigma}_x^{-1} = \begin{bmatrix} \mathbf{M}_0^0 & -\mathbf{M}_0^+ & 0 & \dots & 0 & -\mathbf{M}_0^- \\ -\mathbf{M}_1^- & \mathbf{M}_1^0 & -\mathbf{M}_1^+ & 0 & \dots & 0 \\ \dots & & & & & \dots \\ 0 & \dots & 0 & -\mathbf{M}_{N-1}^- & \mathbf{M}_{N-1}^0 & -\mathbf{M}_{N-1}^+ \\ -\mathbf{M}_N^+ & 0 & \dots & 0 & -\mathbf{M}_N^- & \mathbf{M}_N^0 \end{bmatrix}. \tag{12}$$

If the underlying process is wide–sense stationary, then the matrices $\left\{\mathbf{M}_k^0\right\}$, $\left\{\mathbf{M}_k^+\right\}$, $\left\{\mathbf{M}_k^-\right\}$ do not depend on k and $\mathbf{\Sigma}_x^{-1}$ in (12) is block (tridiagonal and) circulant.

By combining the two characterizations (of reciprocal and Markov process), if one considers the equivalence class of reciprocal processes with dynamics (10), the subclass of Markov processes is such that the blocks in the upper northeast corner and lower southwest corner of $\mathbf{\Sigma}_x^{-1}$ are zero, i.e.

$$\mathbf{M}_N^+ = (\mathbf{M}_0^-)^\top = 0\,.$$

Theorem 3.2 together with the characterization in Theorem 3.3 will be useful in the sequel to provide an alternative derivation in the Gaussian case of the probabilistic graphical model associated to a reciprocal process.

IV. Probabilistic Graphical Models

In this section, we briefly review some relevant theory about probabilistic graphical models needed in the sequel for the derivation of the probabilistic graphical model associated to a reciprocal process (see Section V). We refer the reader to [32], [27], [24], [5] for a thorough treatment of the subject.

Graph–related terminology and background

Let $\mathcal{G} = (V, E)$ be a graph where V denotes the set of vertices and E denotes the set of edges. An edge may be directed or undirected. In case of a directed edge from node i to node j, we say that i is a *parent* of its *child* j. Two nodes i and j are *adjacent* in $\mathcal{G}$ if the directed or undirected edge (i, j) is contained in E. An *undirected path* is a sequence of distinct nodes $\{1, \dots, m\}$ such that there exists a (directed or undirected) edge for each pair of nodes $\{l, l+1\}$ on the path. A graph is *connected* if every pair of points is joined by a path. A graph is *singly-connected* if

there exists only one undirected path between any two nodes in the graph. If this is not the case, the graph is said to be *multiply connected*, or *loopy*. A *(un)directed cycle* is a path such that the beginning and ending nodes on the (un)directed path are the same.

If E contains only undirected edges then the graph $\mathcal{G}$ is an undirected graph (UG). If E contains only directed edges then the graph $\mathcal{G}$ is a directed graph (DG).

Two important classes of graphs for modeling probability distributions that we consider in this paper are UGs and directed acyclic graphs (DAGs), namely directed graphs having no directed cycles.

A graph $\mathcal{G}$ is *complete* if there are edges between all pairs of nodes. A *clique* in an undirected graph is a fully connected set of nodes. A *maximal clique* is a clique that is not a strict subset of another clique. An undirected graph $\mathcal{G}$ is *chordal* if every cycle of length greater than three has an edge connecting nonconsecutive nodes, see e.g. [17]. The *distance* $d(u, v)$ between two vertices u and v in a graph $\mathcal{G}$ is the length of a shortest path between them. If there is no path connecting the two vertices $d(u, v) = \infty$. A shortest path between any two vertices is often called a *geodesic*. The *diameter* of a (connected) graph G, $d(\mathcal{G})$, is the length of any longest geodesic, i.e. $d(\mathcal{G}) = \max\{d(u, v) : u, v \in V\}$.

Probabilistic graphical models

Now that we have introduced some terminology about graphs, we turn to the main object of this paper, namely probabilistic graphical models. *Probabilistic graphical models* are graph–based representations that compactly encode complex distributions over a high-dimensional space. In a probabilistic graphical model, each node represents a random variable and the links express probabilistic relationships between these variables. There are different types of graphical models. Two major classes are *Bayesian networks*, that use directed graphs, and *Markov networks* that are based on undirected graphs. A third class are *factor graphs*.

There are two ways of defining a graphical model: (i) as a representation of a set of independencies, and (ii) as a skeleton for factorizing a distribution. For Bayesian networks, the two definitions are equivalent, while for Markov networks additional assumptions, such as having a positive distribution, are needed to get factorization from independencies (Hammersely–Clifford theorem). The primary definition of Markov networks will thus be in terms of (global) conditional independencies.

The two formalisms, Bayesian and Markov networks, can express different sets of conditional independencies and factorizations, and one or the other may be more appropriate, or even the only suitable, for a particular application. This will be discussed to some extent in the following. In this paper, we will be mainly interested in undirected graphical models, directed ones being mainly useful for expressing causal relationships between random variables.

In this Section, we first briefly introduce Markov and Bayesian networks and then describe relevant theory that allows one to go from a given set of conditional independencies (a distribution) to its graph representation. The graphical model associated to a reciprocal process will be derived in Section V.

Markov Networks

The semantic of undirected graphical models is as follows.

a) Conditional independence property: An undirected graph defines a family of probability distributions which satisfy the following graph separation property.

Property 4.1 (Graph separation property): Let A, B, and C denote three disjoint sets of nodes in an undirected graphical model $\mathcal{H}$ and let us denote by $\mathbf{X}_A$, $\mathbf{X}_B$ and $\mathbf{X}_C$ the corresponding variables in the associated probability distribution P. Then we say that $\mathbf{X}_A \perp\!\!\!\perp \mathbf{X}_B \mid \mathbf{X}_C$ (in P) ($\mathbf{X}_A$ and $\mathbf{X}_B$ are conditionally independent given $\mathbf{X}_C$) whenever (in $\mathcal{H}$) there is no path from a node in A to a node in B which does not pass through a node in C. An alternative way to view this conditional independence test is as follows: remove all nodes in set C from the graph together with any edge that connects to those nodes. If the resulting graph decomposes into multiple connected components such that A and B belong to different components, then $\mathbf{X}_A \perp\!\!\!\perp \mathbf{X}_B \mid \mathbf{X}_C$.

b) Factorization property:

Property 4.2 (Factorization property): Let $\mathcal{H}$ be an undirected graphical model. Let C be a clique and let $\mathbf{X}_C$ be the set of variables in that clique. Let $\mathcal{C}$ denote a set of maximal cliques. Define the following representation of the joint distribution

$$p(\mathbf{x}) = \frac{1}{Z} \prod_{C \in \mathcal{C}} \psi_C(\mathbf{x}_C) \tag{13}$$

where the functions ψ_C can be any nonnegative valued functions (i.e. do not need to sum to 1), and are sometimes referred to as *potential functions* or *compatibility functions* and Z, called

the *partition function*, is a normalization constant chosen in such a way that the probabilities corresponding to all joint assignments sum up to 1. For discrete random variables, it is given by

$$Z = \sum_{\mathbf{x}} \prod_{C \in \mathcal{C}} \psi_C(\mathbf{x}_C) \,.$$

If continuous variables are considered it suffices to replace the summation by an integral.

For positive distributions, the set of distributions that are consistent with the conditional independence statements that can be read from the graph using graph separation and the set of distributions that can be expressed as a factorization of the form (13) with respect to the maximal cliques of the graph are identical. This is the Hammersley–Clifford theorem.

Notice that, differently to what happens for directed graphs, potential functions in undirected graphical models generally do not have a specific probabilistic interpretation as marginal or conditional distributions. Only in special cases, for instance when the undirected graph is constructed by starting with a directed graph, they can admit such interpretation.

Bayesian Networks

A second class of probabilistic graphical models we shall briefly touch upon are *Bayesian networks*. The core of the Bayesian network representation are directed acyclic graphs (DAGs). Similarly to Markov networks, Bayesian networks can be defined both in terms of conditional independencies and factorization properties, but for Bayesian networks, the two definitions are equivalent with no need of additional assumptions. Reading the set of conditional independencies encoded by a Bayesian network again needs testing whether or not the paths connecting two sets of nodes are "blocked", but the definition of "blocked" is this time more involved than it was for undirected graphs. For what concerns the factorization property, in a Bayesian network, factors of the induced distribution represent the conditional distribution of a given variable conditioned on its parents. We do not enter here in further details about Bayesian networks models since, as we shall see, reciprocal processes do not admit a directed graph representation.

Factor graphs

A third type of probabilistic graphical models are factor graphs. A *factor graph* $\mathcal{F}$ is an undirected graph containing two types of nodes: *variable nodes* and *factor nodes*. Suppose we

have a function of several variables $\mathbf{x} = \{\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_N\}$ and that this function factors into a product of several functions, each having some subset of $\{\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_N\}$ as arguments

$$g(\mathbf{x}) = \prod_s f_s(\mathbf{x}_s), \qquad \mathbf{x}_s \text{ subset of variables} \tag{14}$$

This function can be represented by a factor graph having a variable node for each variable $\mathbf{x}_i$, a factor node (depicted by small squares) for each local function f_s and an edge connecting the variable node $\mathbf{x}_i$ to the factor node f_s if and only if $\mathbf{x}_i$ is an argument of f_s.

Notice that every undirected graph can be represented by an equivalent factor graph. The way to do this is to create a factor graph with the same set of variable nodes, and one factor node for each maximal clique in the graph.

From Models to Undirected Graphs

So far, we have been addressing the problem of associating a distribution to a given graphical model via the set of conditional independencies/factorization properties that the graphical model encodes. In this paper, we are interested in finding the probabilistic graphical model associated to a reciprocal process. We are thus interested in the opposite question, and namely: given a process defined by a set of conditional independencies, find a graphical model that encodes such a set, possibly in an "efficient" way. In other words, we want to find a graphical model that encodes all and only the conditional independencies implied by the distribution that we want to represent. Relevant to this aim are the notions of I-map, D-map and P-map, that we are now going to introduce.

Consider a probability distribution P and a graphical model $\mathcal{H}$. Let $CI(P)$ denote the set of conditional independencies satisfied by P and let $CI(\mathcal{H})$ denote the set of all conditional independencies implied by $\mathcal{H}$.

Definition 4.1 (I–map, D–map, P–map): We say that

- $\mathcal{H}$ is an *independence map* (*I–map*) for P if $CI(\mathcal{H}) \subset CI(P)$;
- $\mathcal{H}$ is a *dependence map* (*D–map*) for P if $CI(\mathcal{H}) \supset CI(P)$;
- $\mathcal{H}$ is a *perfect map* (*P-map*) for P if $CI(\mathcal{H}) = CI(P)$.

In other words, if $\mathcal{H}$ is an I–map for P, then every conditional independence statement implied by $\mathcal{H}$ is satisfied by P. If $\mathcal{H}$ is a D–map for P then every conditional independence statement satisfied by P is reflected by $\mathcal{H}$. If it is the case that every conditional independence property of

the distribution is reflected in the graph, and vice versa, then the graph is said to be a perfect map for that distribution. Clearly a fully connected graph will be a trivial I–map for any distribution because it implies no conditional independencies and a graph with no edges will be a trivial D–map for any distribution because it implies every conditional independence.

Generating Minimal I–maps

Back to our original question, we have that an approach to finding a graph that represents a distribution P is simply to take any graph that is an I–map for P. Yet a complete graph is an I-map for any distribution, but there are redundant edges in it. What we are really interested in, are I–maps that represent a family of distributions in a "minimal" way, as specified by the following definition.

Definition 4.2 (Minimal I–map): A *minimal I-map* is an I-map with the property that removing any single edge would cause the graph to no longer be an I-map.

How can we construct a minimal I-map for a distribution P? Here we mention two approaches for constructing a minimal I-map, one based on the pairwise Markov independencies, and the other based on the local independencies (see [32], [33]).

Theorem 4.1: Let P be a positive distribution, and let $\mathcal{H} = (V, E)$ be defined by introducing an edge $(\mathbf{X}, \mathbf{Y})$ for all $\mathbf{X}$, $\mathbf{Y}$ that do *not* satisfy $\mathbf{X} \perp\!\!\!\perp \mathbf{Y} \mid V - \{\mathbf{X}, \mathbf{Y}\}$. Then the Markov network $\mathcal{H}$ is the unique minimal I-map for P.

An alternative approach that uses local independencies, is based on the notion of Markov blanket, that is defined as follows.

Definition 4.3: Consider a graph $\mathcal{H} = (V, E)$. A set U is a *Markov blanket* of $\mathbf{X}$ in a distribution P if $\mathbf{X} \notin U$ and if U is a minimal set of nodes such that

$$(\mathbf{X} \perp\!\!\!\perp V - \{\mathbf{X}\} - U \mid U) \in CI(P)\,. \tag{15}$$

Theorem 4.2: Let P be a positive distribution. For each node $\mathbf{X}$, let $MB_P(\mathbf{X})$ be a minimal set of nodes U satisfying (15). We define a graph $\mathcal{H}$ by introducing an edge $(\mathbf{X}, \mathbf{Y})$ for all $\mathbf{X}$ and $\mathbf{Y} \in MB_P(\mathbf{X})$. Then the Markov network $\mathcal{H}$ is the unique minimal I-map for P.

Nonchordal Markov networks do not admit a Bayesian network as a perfect map

It turns out that some distributions can be perfectly represented by a directed graphical model while others can be perfectly represented by an undirected one. On the other hand, some sets

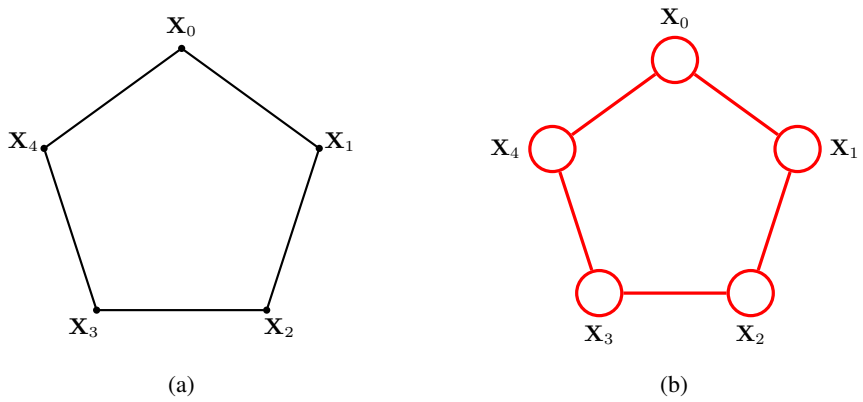

Fig. 1: Wrapped time line (on the left) and probabilistic graphical model (on the right) for a reciprocal process on $\mathcal{I} = [0, 4]$.

of independence assumptions can be perfectly represented both by a Bayesian network and by a Markov network. This is the case of undirected chordal graphs. The precise statement is as follows.

Theorem 4.3: Let $\mathcal{H}$ be a Markov network. Then there is a Bayesian network $\mathcal{G}$ such that $CI(\mathcal{H}) = CI(\mathcal{G})$ if and only if $\mathcal{H}$ is chordal.

V. Probabilistic Graphical Models of Reciprocal Processes

We are now ready to state our main result, namely to find the graphical model associated to a reciprocal process. We first derive the minimal I–map associated to a reciprocal process (Theorem 5.1) and then show that this minimal I–map is indeed also a P-map (perfect map) for the reciprocal process (Theorem 5.2).

Theorem 5.1: The undirected graphical model composed of the $N + 1$ nodes $\mathbf{X}_0, \mathbf{X}_1, \ldots, \mathbf{X}_N$ arranged in a loop (see Figure 1b) is the unique minimal I–map for a reciprocal process on $\mathcal{I} = [0, N]$ with cyclic boundary conditions.

Proof via Theorem 4.2: Let P_R denote the distribution of a reciprocal process on $\mathcal{I}$. A Markov blanket of $\mathbf{X}_k$ in P_R is the set $U_R = \{\mathbf{X}_{k-1}, \mathbf{X}_{k+1}\}$ (where for $k = 0$ and $k = N$, $k \pm 1$ has to be read modulo $N + 1$). The undirected graphical model in Figure 1b thus follows by using the construction criterion in Theorem 4.2. ■

For Gaussian reciprocal processes, one can also exploit the characterization in terms of sparsity pattern of the inverse covariance matrix of Theorem 3.3 and proceed as follows.

Proof via Theorem 4.1: By Theorem 3.2, setting the (i, j)–th element of the inverse covariance matrix to zero has the probabilistic interpretation that the i–th and j–th components of the underlying Gaussian random vector are conditionally independent given the other components. The undirected graphical model in Figure 1b thus follows by the characterization of reciprocal processes in Theorem 3.3, using the construction criterion in Theorem 4.1. ■

Theorem 5.2: The undirected graphical model composed of the $N+1$ nodes $\mathbf{X}_0, \mathbf{X}_1, \ldots, \mathbf{X}_N$ arranged in a loop (see Figure 1b) is a P–map for a reciprocal process on $\mathcal{I} = [0, N]$ with cyclic boundary conditions.

Proof: Consider a reciprocal process on the interval $\mathcal{I} = [0, N]$ with cyclic boundary conditions $\mathbf{X}_{-1} = \mathbf{X}_N$, $\mathbf{X}_{N+1} = \mathbf{X}_0$. Because cyclic boundary conditions hold, one may think to the process as defined on the wrapped timeline in Figure 1a. The thesis follows from the definition of reciprocal process and the Separation Property 4.1 by noting that the extremes of each interval on the (wrapped) timeline define a set (pair) of nodes on the corresponding graphical model that decomposes the graph into multiple connected components such that nodes corresponding to the "interior" and the "exterior" of the interval belong to different components. ■

Reciprocal processes do not admit a directed graph as perfect map

The Markov network in Figure 1b is *not* chordal, thus, by Theorem 4.3 reciprocal processes do *not* admit a directed graph as perfect map. This is in contrast with Markov processes that admit both a directed and an undirected graphical model as perfect map.

Factor graph representation of a Reciprocal Process

As observed above, every undirected graph can be represented by an equivalent factor graph having the same set of variable nodes, and one factor node for each maximal clique in the graph. The factor graph corresponding to the undirected graph in Figure 1b is shown in Figure 2.

Link with the four nodes single loop undirected graphical model by Pearl

The single loop undirected graphical model in Figure 1b has been considered in [32, p. 90] (see also [24]), where it has been used to model the spread of a disease or of a misconception among individuals who only engage in pairwise activities, and is used as a motivating example for the

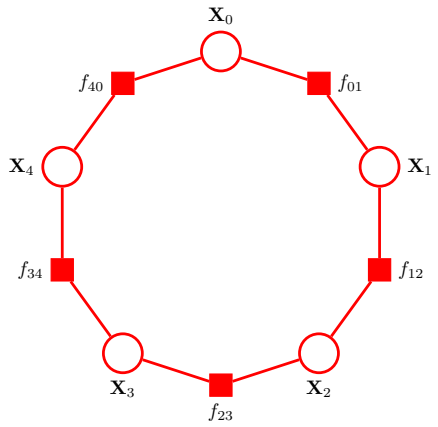

Fig. 2: Factor Graph associated with a reciprocal process on $\mathcal{I} = [0, 4]$.

introduction of Markov networks, since, as observed above, the underlying set of conditional independencies does not admit a Bayesian network as a perfect map. Nevertheless, to the best of our knowledge, this is the first time that such graphical model is associated to a reciprocal process (distribution). This fills a gap in the Graphical Models literature, by bridging a well–known graphical structure with the class of reciprocal processes studied in the Statistics and in the Control communities. Moreover it opens the way to new applications of reciprocal processes, e.g. to the study of the spread of certain diseases or in opinion formation in social networks, that do not seem to have been explored so far in the literature.

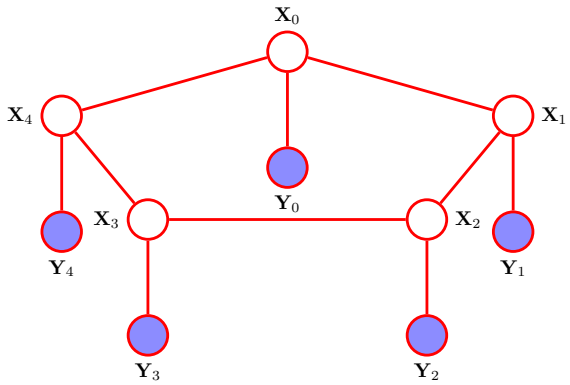

Fig. 3: Probabilistic graphical model associated to a hidden reciprocal model on $\mathcal{I} = [0, 4]$.

VI. Smoothing of Reciprocal Processes via Belief Propagation

Consider a reciprocal process $\{\mathbf{X}_k\}$ and a second process $\{\mathbf{Y}_k\}$, where, given the state sequence $\{\mathbf{X}_k\}$, the $\{\mathbf{Y}_k\}$ are independent random variables, and, for all k, the conditional probability distribution of $\mathbf{Y}_k$ depends only on $\mathbf{X}_k$. In applications, $\{\mathbf{X}_k\}$ represents a "hidden" process which is not directly observable, while the observable process $\{\mathbf{Y}_k\}$ represents "noisy observations" of the hidden process. We shall refer to the pair $\{\mathbf{X}_k, \mathbf{Y}_k\}$ as a *hidden reciprocal model*. The corresponding probabilistic graphical model is illustrated in Figure 3. The (fixed–interval) *smoothing problem* is to compute, for all $k \in [0, N]$, the conditional distribution of $\mathbf{X}_k$ given $\mathbf{Y}_0, \dots, \mathbf{Y}_N$. One of the most widespread algorithms for performing inference (solving the smoothing problem) in the graphical models literature is the *belief propagation* algorithm [32], [24], [5]. This is reviewed in Section VI-A and specialized for reciprocal processes in Section VI-B. Notice that this approach is distribution independent and holds indeed both for continuous and discrete–valued random variables/reciprocal process (even if in the Gaussian case, it may result convenient to rewrite the iteration, which lives in the infinite dimensional space of nonnegative measurable functions, in the finite dimensional spaces of mean vectors and covariance matrices, see [7] for a particularization to reciprocal processes). In this Section, we state the algorithm for continuous–valued variables, the discrete variables case following immediately by replacing integrals with summations where appropriate.

A. Belief Propagation (a.k.a. sum–product) algorithm

Let $\mathcal{H} = (E, V)$ be an undirected graphical model over the variables $\{\mathbf{X}_0, \dots, \mathbf{X}_N\}$, $\mathbf{X}_i \in \mathcal{X}$, $i = 0, \dots, N$. In Section IV, we have seen that the joint distribution associated with $\mathcal{H}$ can be factored as

$$p(\mathbf{x}) = \frac{1}{Z} \prod_{C \in \mathcal{C}} \psi_C(\mathbf{x}_C)\,, \tag{16}$$

where $\mathcal{C}$ denotes a set of maximal cliques in the graph. In the following, we will be interested in pairwise Markov random fields – i.e. a Markov random field in which the joint probability factorizes into a product of bivariate potentials (potentials involving only two variables) – where each unobserved node $\mathbf{X}_i$ has an associated observed node $\mathbf{Y}_i$. Factorization (16) then becomes

$$p(\mathbf{x}_{0:N}, \mathbf{y}_{0:N}) = \prod_{(i,j) \in E} \psi_{ij}(\mathbf{x}_i, \mathbf{x}_j) \prod_i \psi_i(\mathbf{x}_i, \mathbf{y}_i)\,, \tag{17}$$

where the $\psi_{ij}(\mathbf{x}_i, \mathbf{x}_j)$'s are often referred to as the *edge potentials* and the $\psi_i(\mathbf{x}_i, \mathbf{y}_i)$'s are often referred to as the *node potentials*. The problem we are interested in is finding posterior marginals of the type $p(\mathbf{x}_i \mid \mathbf{y}_{0:N})$ for some hidden variable $\mathbf{X}_i$. The basic idea behind belief propagation is to exploit the factorization properties of the distribution to allow efficient computation of the marginals. Indeed, since the scope of the factors in (17) is limited, this allows us to "push in" some of the integrals, performing them over a subset of variables at a time. To fix ideas, consider the graph in Figure 4 and suppose we want to compute the conditional marginal $p(\mathbf{x}_0 \mid \mathbf{y}_{0:3})$. A naive application of the definition, would suggest that $p(\mathbf{x}_0 \mid \mathbf{y}_{0:3})$ can be obtained by integrating the joint distribution over all variables except $\mathbf{X}_0$ and then normalize

$$p(\mathbf{x}_0 \mid \mathbf{y}_{0:3}) \propto \int_{\mathbf{x}_1} \int_{\mathbf{x}_2} \int_{\mathbf{x}_3} p(\mathbf{x}, \mathbf{y}) d\mathbf{x}_1 d\mathbf{x}_2 d\mathbf{x}_3 \,. \tag{18}$$

Nevertheless notice that the joint distribution can be factored as:

$$\begin{aligned} p(\mathbf{x}_{0:3}, \mathbf{y}_{0:3}) = \psi_0(\mathbf{x}_0)\psi_{01}(\mathbf{x}_0, \mathbf{x}_1)\psi_1(\mathbf{x}_1)\psi_{12}(\mathbf{x}_1, \mathbf{x}_2) \\ \psi_2(\mathbf{x}_2)\psi_{13}(\mathbf{x}_1, \mathbf{x}_3)\psi_3(\mathbf{x}_3) \,. \end{aligned} \tag{19}$$

By plugging in factorization (19) into equation (18) and interchanging the integrals and products order, we obtain

$$\begin{aligned} p(\mathbf{x}_0 \mid \mathbf{y}_{0:3}) \propto \psi_0(\mathbf{x}_0) \Bigg[\int_{\mathbf{x}_1} \psi_{01}(\mathbf{x}_0, \mathbf{x}_1)\psi_1(\mathbf{x}_1) \int_{\mathbf{x}_2} \psi_{12}(\mathbf{x}_1, \mathbf{x}_2)\psi_2(\mathbf{x}_2) \\ \int_{\mathbf{x}_3} \psi_{13}(\mathbf{x}_1, \mathbf{x}_3)\psi_3(\mathbf{x}_3) \Bigg] \end{aligned} \tag{20}$$

This simple operation forms the basis of the belief propagation algorithm and it can be given an interpretation in terms of passing of local messages around the graph. Most importantly, notice that it allows to reduce the computational cost of the computation of the posterior marginal from exponential to linear in the number of nodes. Indeed for a tree structured graph with random variables taking values in a finite alphabet $\mathcal{X}$, the computational cost passes from $O(|\mathcal{X}|^N)$ (computational cost of the brute force marginalization in (18)) to $O(N|\mathcal{X}|^2)$ (computational cost of the "principled" marginalization in (20)), where $|\mathcal{X}|$ denotes the cardinality of the set $\mathcal{X}$. In its general form, the belief propagation algorithm reads as follows.

Algorithm 1 (Belief propagation): Let $\mathbf{X}_i$ and $\mathbf{X}_j$ be two neighboring nodes in the graph. We denote by m_{ij} the message that node $\mathbf{X}_i$ sends to node $\mathbf{X}_j$, by m_{ii} the message that $\mathbf{Y}_i$ sends

to $\mathbf{X}_i$, and by b_i the "belief" (estimated posterior marginal) at node $\mathbf{X}_i$. The belief propagation algorithm is as follows:

$$m_{ij}(\mathbf{x}_j) = \alpha \int_{\mathbf{x}_i} \psi_{ij}(\mathbf{x}_i, \mathbf{x}_j) m_{ii}(\mathbf{x}_i) \prod_{k \in \partial i \setminus j} m_{ki}(\mathbf{x}_i) \tag{21a}$$

$$b_i(\mathbf{x}_i) = \beta \; m_{ii}(\mathbf{x}_i) \prod_{k \in \partial i} m_{ki}(\mathbf{x}_i) \tag{21b}$$

where ∂i denotes the set of neighbors of node $\mathbf{X}_i$ and α and β are normalization constants so that messages and beliefs integrate to one.

For example, if one considers (20), setting $m_{ii}(\mathbf{x}_i) := \psi_i(\mathbf{x}_i)$ and applying definition (21a) for the messages, taking into account that $\psi_{ij}(\mathbf{x}_i, \mathbf{x}_j) = \psi_{ji}(\mathbf{x}_j, \mathbf{x}_i)$, (20) becomes

$$\begin{aligned} p(\mathbf{x}_0 | \, \mathbf{y}_{0:3}) &\propto m_{00}(\mathbf{x}_0) \left\{ \int_{\mathbf{x}_1} \psi_{01}(\mathbf{x}_0, \mathbf{x}_1) m_{11}(\mathbf{x}_1) m_{21}(\mathbf{x}_1) m_{31}(\mathbf{x}_1) \right\} \\ &= m_{00}(\mathbf{x}_0) \cdot m_{10}(\mathbf{x}_0) \end{aligned}$$

which is of the form (21b), where the posterior marginal $p(\mathbf{x}_0 \mid \mathbf{y}_{0:3})$ is computed as the product of incoming messages at node $\mathbf{X}_0$.

Before going on, some observations about the belief propagation algorithm are in order. Observed nodes do not receive messages, and they always transmit the same message m_{ii}. The normalization of messages in equation (21a) is not theoretically necessary (whether the messages are normalized or not, the beliefs b_i will be identical) but helps improving numerical stability of the algorithm. Equation (21a) does not specify the order in which the messages are updated. While a sequential scheduling policy is possible in tree–structured graphs, with messages sequentially computed starting from leaf nodes once incoming messages at a given node become available, this is not viable in a loopy network. In this paper, following [38], we assume that all nodes simultaneously update their messages in parallel. This naturally leads to *loopy belief propagation*, where the update rule (21a) is applied to graphs that are not a tree. This is the case of reciprocal processes, that we are going to treat in the next section.

B. Belief Propagation for Hidden Reciprocal Models

If the considered graph is the single–loop hidden reciprocal model in Figure 3, expressions (21a) and (21b) for the message and belief updates simplify, each node having only two neighbors. Moreover we can distinguish between two classes of messages, one propagating in the

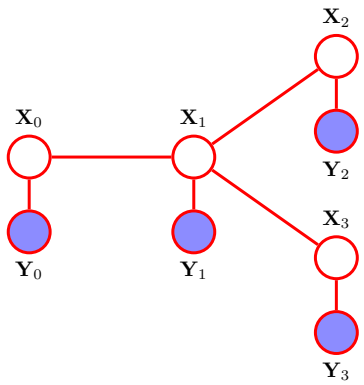

Fig. 4: A graphical model with four unobserved nodes $\mathbf{X}_0, \ldots, \mathbf{X}_3$ and four observed nodes $\mathbf{Y}_0, \ldots, \mathbf{Y}_3$.

direction of increasing indexes (clockwise) and one propagating in the direction of decreasing indexes (anticlockwise) in the loop. The overall algorithm with parallel scheduling policy is as follows:

Algorithm 2 ((Parallel) loopy belief propagation algorithm for reciprocal processes): 1) Initialize all messages $\mathbf{m}_{ks}^{(0)}$ to some initial value $\bar{\mathbf{m}}_{ks}^{(0)}$.

2) For $t \in \{0, 1, \ldots, t_{max}\}$, for all $k \in \{0, 1, \ldots, N\}$

$$m_{k-1,k}^{(t+1)}(\mathbf{x}_k) = \alpha_f \int_{\mathbf{x}_{k-1}} \psi_{k-1,k}(\mathbf{x}_{k-1}, \mathbf{x}_k) m_{k-1,k-1}(\mathbf{x}_{k-1}) m_{k-2,k-1}^{(t)}(\mathbf{x}_{k-1}) \tag{22a}$$

$$m_{k+1,k}^{(t+1)}(\mathbf{x}_k) = \alpha_b \int_{\mathbf{x}_{k+1}} \psi_{k+1,k}(\mathbf{x}_{k+1}, \mathbf{x}_k) m_{k+1,k+1}(\mathbf{x}_{k+1}) m_{k+2,k+1}^{(t)}(\mathbf{x}_{k+1}) \,. \tag{22b}$$

3) For each $\mathbf{X}_k$, $k = 0, \ldots, N$, compute the posterior marginals

$$b_k(\mathbf{x}_k) = \beta \;\; m_{kk}(\mathbf{x}_k) \left[m_{k-1,k}^{(t_{max}+1)}(\mathbf{x}_k) \cdot m_{k+1,k}^{(t_{max}+1)}(\mathbf{x}_k) \right] . \tag{23}$$

For tree-structured graphs, when t_{max} is larger than the diameter of the tree (the length of the longest shortest path between any two vertices of the graph), the algorithm converges to the correct marginal. This is not the case for reciprocal processes, whose associated graph is the single loop network in Figure 3. Convergence of the iteration for a hidden reciprocal model will be discussed in Section IX. The argument we will use is based on contraction properties of positive operators with respect to the Hilbert metric, that we are now going to introduce.

VII. Hilbert metric

The Hilbert metric was introduced in [18] and is defined as follows. Let $\mathcal{B}$ be a real Banach space and let $\mathcal{K}$ be a closed solid cone in $\mathcal{B}$, that is a closed subset $\mathcal{K}$ with the properties that (i) $\lambda\mathcal{K} \subset \mathcal{K}$ for all $\lambda \geq 0$ (ii) the interior of $\mathcal{K}$, $\mathcal{K}^+$, is non–empty; (iii) $\mathcal{K} + \mathcal{K} \subseteq \mathcal{K}$; (iv) $\mathcal{K} \cap -\mathcal{K} = \{0\}$. Define the partial order

$$x \preceq y \Leftrightarrow y - x \in \mathcal{K}\,,$$

and for $x, y \in \mathcal{K}\backslash\{0\}$, let

$$M(x,y) := \inf\{\lambda | x - \lambda y \preceq 0\}$$

$$m(x,y) := \sup\{\lambda | x - \lambda y \succeq 0\}$$

The Hilbert metric $d_{\mathcal{H}}(\cdot,\cdot)$ induced by $\mathcal{K}$ is defined by

$$d_{\mathcal{H}}(x,y) := \log\left(\frac{M(x,y)}{m(x,y)}\right), \quad x, y \in \mathcal{K}\backslash\{0\}\,. \tag{24}$$

For example, if $\mathcal{B} = \mathbb{R}^n$ and the cone $\mathcal{K}$ is the positive orthant, $\mathcal{K} = \mathcal{O} := \{(x_1, \ldots, x_n) : x_i \geq 0,\ 1 \leq i \leq n\}$, then $M(\mathbf{x},\mathbf{y}) = \max_i(x_i/y_j)$ and $m(\mathbf{x},\mathbf{y}) = \min_i(x_i/y_i)$ and the Hilbert metric can be expressed as

$$d_{\mathcal{H}}(\mathbf{x},\mathbf{y}) = \log\frac{\max_i(x_i/y_i)}{\min_i(x_i/y_i)}\,.$$

On the other hand, if $\mathcal{B} = \mathcal{S} := \{\mathbf{X} = \mathbf{X}^\top \in \mathbb{R}^{n\times n}\}$ is the set of symmetric matrices and $\mathcal{K} = \mathcal{P} := \{\mathbf{X} \geq 0 \mid \mathbf{X} \in \mathcal{S}\}$ is the cone of positive semidefinite matrices, then for $\mathbf{X}, \mathbf{Y} > 0$, $M(\mathbf{X},\mathbf{Y}) = \lambda_{max}\left(\mathbf{X}\mathbf{Y}^{-1}\right)$ and $m(\mathbf{X},\mathbf{Y}) = \lambda_{min}\left(\mathbf{X}\mathbf{Y}^{-1}\right)$. Hence the Hilbert metric is

$$d_{\mathcal{H}}(\mathbf{X},\mathbf{Y}) = \log\frac{\lambda_{max}\left(\mathbf{X}\mathbf{Y}^{-1}\right)}{\lambda_{min}\left(\mathbf{X}\mathbf{Y}^{-1}\right)}\,.$$

An important property of the Hilbert metric is the following. The Hilbert metric is a *projective metric* on $\mathcal{K}$ i.e. it is nonnegative, symmetric, it satisfies the triangle inequality and is such that, for every $x, y \in \mathcal{K}$, $d_{\mathcal{H}}(x,y) = 0$ if and only if $x = \lambda y$ for some $\lambda > 0$. It follows easily that $d_{\mathcal{H}}(x,y)$ is constant on rays, that is

$$d_{\mathcal{H}}(\lambda x, \mu y) = d_{\mathcal{H}}(x,y) \quad \text{for } \lambda, \mu > 0\,. \tag{25}$$

A second relevant property is in connection with positive operators. In [4] (see also [6]) it has been shown that linear positive operators contract the Hilbert metric. This can be used to provide a geometric proof of the Perron–Frobenius theorem and, in turn, to prove attractiveness properties of linear positive systems. This is the subject of the next section.

VIII. Positive systems

A linear time invariant system $\mathbf{x}_{k+1} = \mathbf{A}\mathbf{x}_k$ over the positive orthant O is positive if the mapping $\mathbf{A}$ takes O into itself. Positive systems have a long history in the literature, both because of the relevance of the property for applications (the positivity constraint arises quite naturally when modeling real systems whose state variables represent quantities that are intrinsically nonnegative, such as pressures, concentrations, population levels, etc) and because the property significantly restricts the behavior, as established by Perron–Frobenius theory: if the cone invariance is strict, that is, if the boundary of the cone is eventually mapped to the interior of the cone, then the asymptotic behavior of the system lies on a one dimensional object. In Section VIII-A, we briefly review contraction properties of positive linear operators as derived by Birkhoff [4] (see also [6]) and then show how they can be used to prove existence of a fixed point for a linear time invariant positive dynamical system which is also a global attractor (Section VIII-B).

A. Positive linear operators contract the Hilbert metric

Let (X, d) be a metric space. We recall that a mapping $f : X \to X$ is called a *contraction* with respect to d if there exists $0 \le K < 1$ such that

$$d(f(x), f(y)) \le K d(x, y), \quad \text{for all } x, y \in X. \tag{26}$$

A map A from $\mathcal{B}$ to $\mathcal{B}$ is said to be *non–negative* if it takes $\mathcal{K}$ into itself, i.e.

$$A : \mathcal{K} \to \mathcal{K},$$

and *positive* if it takes the interior of $\mathcal{K}$ into itself, i.e.

$$A : \mathcal{K}^+ \to \mathcal{K}^+.$$

For a positive map define its *contraction ratio*

$$k(A) := \inf\{\lambda : d(Ax, Ay) \le \lambda d(x, y)\ \forall x, y \in \mathcal{K}^+\} \tag{27}$$

and *projective diameter*

$$\Delta(A) := \sup\{d(Ax, Ay) : x, y \in \mathcal{K}^+\}. \tag{28}$$

It is easy to show that a non–negative linear map does not expand the Hilbert metric [23]. In [4] (see also [6]), Birkhoff showed that positivity of a linear mapping implies contraction in the

Hilbert metric, a result that paved the way to many contraction–based results in the literature of positive operators. The formal statement is as follows.

Theorem 8.1: If $x, y \in \mathcal{K}$, then the following holds

(i) if A is a non–negative linear map on $\mathcal{K}$, then $d_{\mathcal{H}}(Ax, Ay) \leq d_{\mathcal{H}}(x, y)$, i.e. the Hilbert metric contracts weakly under the action of a non–negative linear transformation.

(ii) [Birkhoff, 1957] If A is a positive linear map in $\mathcal{B}$, then

$$k(A) = \tanh \frac{1}{4} \Delta(\mathrm{A}) \,. \tag{29}$$

In other words, if the diameter $\Delta(A)$ is finite, then positivity of a mapping implies strict contraction of the Hilbert metric.

In the next section we will see how contraction properties of positive linear maps with respect to the Hilbert metric can be used to explain the asymptotic behavior of a positive time–invariant system.

B. Asymptotic Behavior of Positive Dynamical Systems via Contraction of the Hilbert metric

Exploiting contraction properties of positive maps, in [4] (see also [6]) Birkhoff provided an alternative proof of the the Perron–Frobenius theorem as a special case of the Banach fixed-point theorem. With respect to others fixed-point arguments used to prove the Perron-Frobenius theorem (see e.g. [14]), this proof has the advantage that not only it yields to the existence of a positive eigenvector $\mathbf{x}_f$, but also to convergence to this same eigenvector (for the latter, in the approach in [14], one still needs to show how positivity implies that the eigenvalue associated with $\mathbf{x}_f$ dominates all the other eigenvalues). Along the same lines, we exploit contraction properties of positive maps with respect to the Hilbert metric to prove existence of a fixed point of the projective space for a linear time–invariant positive dynamical system which is a global attractor for the system. This will be used in the next section to prove convergence of loopy belief propagation for reciprocal processes, where we will show that the underlying iteration is indeed a positive system.

Theorem 8.2: Consider the dynamical system $\mathbf{x}_{k+1} = \mathbf{A}\mathbf{x}_k$ with $\mathbf{A}$ a $D \times D$ matrix with non–negative entries and suppose that $\mathbf{A}$ is such that there exists an integer h such that

$$[\mathbf{A}^h]_{r,s} > 0, \quad \forall\, r,\ s \in \{1, 2, \ldots, D\} \,,$$

(i.e. the matrix $\mathbf{A}$ is primitive). Then there exists a unique positive eigenvector $\mathbf{x}_f \in \mathcal{O}^+$ such that for all non–negative $\mathbf{x}_0 \in \mathcal{O}\backslash\{0\}$, $\mathbf{A}^n \mathbf{x}_0$ converges in direction to $\mathbf{x}_f$, i.e.

$$d_{\mathcal{H}}(\mathbf{A}^n \mathbf{x}_0, \mathbf{x}_f) \to 0 \quad \text{as } n \to \infty$$

and the rate of convergence is at least linear (i.e. the error decreases exponentially).
To prove Theorem 8.2 we need the two following lemmas.

Lemma 8.1: [6] Consider the cone $\mathcal{K} = \mathcal{O} := \{(x_1, \ldots, x_D) : x_i \geq 0,\ 1 \leq i \leq D\}$ (positive orthant) and let U denote the unit sphere in $\mathcal{B} = \mathbb{R}^D$. Then the metric space $E := \{\mathcal{O}^+ \cap U, d_{\mathcal{H}}\}$ is complete.

Lemma 8.2: [4] Let $\mathbf{A} = (a_{ij})$ be a $D \times D$ matrix with $a_{ij} > 0$, for all i, j. Then $\mathbf{A}$ is a positive map with finite projective diameter given by

$$\Delta(\mathbf{A}) = \max\left\{\log \frac{a_{ij}a_{pq}}{a_{iq}a_{pj}} : 1 \leq i,\ j,\ p,\ q \leq D\right\} < \infty\,.$$

Proof of Theorem 8.2: $\mathbf{A}^h$ is a positive linear mapping in the interior of the positive orthant $\mathcal{O}^+$ in $\mathbb{R}^D$ with finite projective diameter (see Lemma 8.2). Then $F(\mathbf{x}) := \frac{\mathbf{A}^h\mathbf{x}}{\|\mathbf{A}^h\mathbf{x}\|}$ is a map from $E := \{\mathcal{O}^+ \cap U, d_{\mathcal{H}}\}$ into E and is the composition of a strict contraction (see Theorem 8.1(ii)) and a normalizing isometry (see (25)). Since the metric space E is complete (Lemma 8.1), then, by the Banach fixed-point theorem, there exists a unique fixed point $\bar{\mathbf{x}}_f$ in E such that $F\left(\bar{\mathbf{x}}_f\right) = \bar{\mathbf{x}}_f$, i.e. $\bar{\mathbf{x}}_f$ is a strictly positive eigenvector of $\mathbf{A}^h$ (and of $\mathbf{A}$), associated to a positive eigenvalue, that, starting from an arbitrary $\bar{\mathbf{x}}_0 \in E$ (indeed for every $\bar{\mathbf{x}}_0 \in \mathcal{O} \cap U$) can be computed as the limit of the sequence $\bar{\mathbf{x}}_k = F(\bar{\mathbf{x}}_{k-1})$ so that for every $\epsilon > 0$, there exists a natural number $\bar{k}$ such that, for all $k > \bar{k}$, $d_{\mathcal{H}}(\bar{\mathbf{x}}_k, \bar{\mathbf{x}}_f) = d_{\mathcal{H}}(F^k(\bar{\mathbf{x}}_0), \bar{\mathbf{x}}_f) = d_{\mathcal{H}}\left(\frac{\mathbf{A}^{kh}\bar{\mathbf{x}}_0}{\|\mathbf{A}^{kh}\bar{\mathbf{x}}_0\|}, \bar{\mathbf{x}}_f\right) = d_{\mathcal{H}}(\mathbf{A}^{kh}\bar{\mathbf{x}}_0, \bar{\mathbf{x}}_f) = d_{\mathcal{H}}(\mathbf{A}^{kh}\mathbf{x}_0, \mathbf{x}_f) < \epsilon$. That the rate of convergence is at least linear follows by the fact that $\mathbf{A}^n$ is a contractive map. QED. ■

We are now ready to state our third contribution, namely to revisit convergence analysis of loopy belief propagation for a single loop network via contraction of the Hilbert metric. This is the subject of the next Section.

IX. Convergence of Loopy Belief Propagation for Reciprocal Processes

When the graph is singly connected, local propagation rules are guaranteed to converge to the correct posterior probabilities [24]. For general graphs with loops, theoretical understanding of

the performance of local propagation schemes is an active field of research (see [38], [39], [19], [31], [37], [30] and references therein). Convergence of loopy belief propagation for networks with a single loop has been studied in [38] where it has been shown that, for latent variables taking values in a finite alphabet, the estimated beliefs converge and accuracy of the approximation has been analyzed. A third contribution of this paper is to provide an alternative argument to explain convergence of loopy belief propagation that leverages on contraction properties of positive operators with respect to the Hilbert metric. This argument is geometric in nature and as such can be generalized, e.g. to the Gaussian case, where convergence of the iteration, which is this time a nonlinear map on the cone of positive semidefinite matrices, can again be traced back to contraction properties of the Hilbert metric (see [7] for details). The Section is organized as follows. In Section IX-A, we show that, for a single loop network, convergence analysis of belief propagation essentially boils down to the analysis of the asymptotic behavior of a linear positive dynamical system. Leveraging on results in Section VIII, convergence of the update is then derived as a consequence of contraction properties of positive linear operators with respect to the Hilbert metric (Section IX-B). Necessary and sufficient conditions for convergence of loopy belief propagation borrowed from the theory of linear positive systems are introduced in Section IX-C. Following [38], accuracy of the approximated posterior is discussed in Section IX-D.

A. Belief propagation and positive systems

In case of finite state space, the belief propagation algorithm can be written in matrix notation as follows. If one denotes with $\mathbf{M}_{ij}$ the transition matrix associated with the edge potential $\psi_{ij}(\mathbf{x}_i, \mathbf{x}_j)$ and by $\mathbf{m}_{ij}$ (resp., $\mathbf{m}_{ii}$) the vector messages obtained by staking the $m_{ij}(\mathbf{x}_i)$'s (resp., the $m_{ii}(\mathbf{x}_i)$'s) for each value of $\mathbf{X}_i$ in $\mathcal{X} = \{0, \dots, D-1\}$, namely

$$\mathbf{m}_{ij} = \begin{bmatrix} m_{ij}(0) \\ \vdots \\ m_{ij}(D-1) \end{bmatrix}, \qquad \mathbf{m}_{ii} = \begin{bmatrix} \psi_i(0) \\ \vdots \\ \psi_i(D-1) \end{bmatrix},$$

terms of the form $\sum_{x_i} \psi_{ij}(\mathbf{x}_i, \mathbf{x}_j) m_{ki}(\mathbf{x}_i)$ can be expressed in matrix notation as $\mathbf{M}_{ij}\mathbf{m}_{ki}$. Moreover we denote

$$\mathbf{b}_i := \begin{bmatrix} b_i(0) \\ \vdots \\ b_i(D-1) \end{bmatrix},$$

and indicate by $\odot$ the Hadamard (entrywise) product between two vectors of the same size. Thus, for a reciprocal chain, the message updates (22a)–(22b) can be expressed in matrix notation as

$$\mathbf{m}_{k-1,k}^{(t+1)} = \alpha_f \ \mathbf{M}_{k-1,k} \left(\mathbf{m}_{k-1,k-1} \odot \mathbf{m}_{k-2,k-1}^{(t)} \right) \tag{30}$$

$$\mathbf{m}_{k+1,k}^{(t+1)} = \alpha_b \ \mathbf{M}_{k+1,k} \left(\mathbf{m}_{k+1,k+1} \odot \mathbf{m}_{k+2,k+1}^{(t)} \right) \tag{31}$$

and the beliefs (23) as

$$\mathbf{b}_k = \beta \ \left(\mathbf{m}_{kk} \odot \mathbf{m}_{k-1,k}^{(t_{max}+1)} \odot \mathbf{m}_{k+1,k}^{(t_{max}+1)} \right). \tag{32}$$

Now consider the hidden reciprocal model in Figure 3. Without loss of generality, consider the belief at node $\mathbf{X}_0$ at a certain time $t + N + 1$, which is given by

$$\mathbf{b}_0^{(t+N+1)} = \beta \ \mathbf{m}_{00} \odot \left(\mathbf{m}_{N0}^{(t+N+1)} \odot \mathbf{m}_{10}^{(t+N+1)} \right). \tag{33}$$

By the message update equation (30), the message that $\mathbf{X}_N$ sends to $\mathbf{X}_0$ at time $t+N+1$ depends on the message that $\mathbf{X}_N$ received from $\mathbf{X}_{N-1}$ at time $t + N$

$$\mathbf{m}_{N0}^{(t+N+1)} = \alpha_f \mathbf{M}_{N0} \left(\mathbf{m}_{NN} \odot \mathbf{m}_{N-1,N}^{(t+N)} \right). \tag{34}$$

Similarly, the message that $\mathbf{X}_{N-1}$ sends to $\mathbf{X}_N$ at time $t + N$ depends on the message that $\mathbf{X}_{N-1}$ received from $\mathbf{X}_{N-2}$ at time $t + N - 1$

$$\mathbf{m}_{N-1,N}^{(t+N)} = \alpha_f \mathbf{M}_{N-1,N} \left(\mathbf{m}_{N-1,N-1} \odot \mathbf{m}_{N-2,N-1}^{(t+N-1)} \right) \tag{35}$$

and so on. One can continue expressing each message in terms of the one received from the neighbor until we go back in the loop to $\mathbf{X}_0$: the message that $\mathbf{X}_0$ sends to $\mathbf{X}_1$ is a function of the message that $\mathbf{X}_N$ sent to $\mathbf{X}_0$

$$\mathbf{m}_{01}^{(t+1)} = \alpha_f \mathbf{M}_{01} \left(\mathbf{m}_{00} \odot \mathbf{m}_{N0}^{(t)} \right). \tag{36}$$

By putting together (34)–(36), one gets that the message that $\mathbf{X}_N$ sends to $\mathbf{X}_0$ at a given time step depends on the message that $\mathbf{X}_N$ sent to $\mathbf{X}_0$ $N+1$ time steps ago. In particular, if we denote by $\mathbf{C}_{N0}$ the matrix

$$\mathbf{C}_{N0} = \mathbf{M}_{N0}\mathbf{D}_N\mathbf{M}_{N-1,N}\mathbf{D}_{N-1} \cdot \cdots \cdot \mathbf{M}_{01}\mathbf{D}_0 \tag{37}$$

where the $\mathbf{D}_i$'s are the diagonal matrices whose diagonal elements are the entries of the constant messages $\mathbf{m}_{ii}$, the message that $\mathbf{X}_N$ sends to $\mathbf{X}_0$ satisfy the recursion

$$\mathbf{m}_{N0}^{(t+N+1)} = \alpha_f \mathbf{C}_{N0} \mathbf{m}_{N0}^{(t)} . \tag{38}$$

In a similar way, we can express the message that $\mathbf{X}_1$ sends to $\mathbf{X}_0$ at a given time step as a function of the message that $\mathbf{X}_1$ sent to $\mathbf{X}_0$ $N+1$ time steps ago

$$\mathbf{m}_{10}^{(t+N+1)} = \alpha_b \mathbf{C}_{10} \mathbf{m}_{10}^{(t)} , \tag{39}$$

where the matrix $\mathbf{C}_{10}$ is given by

$$\mathbf{C}_{10} = \mathbf{M}_{10}\mathbf{D}_1\mathbf{M}_{21}\mathbf{D}_2 \cdot \cdots \cdot \mathbf{M}_{0N}\mathbf{D}_0 . \tag{40}$$

Furthermore, since for any two nodes $\mathbf{X}_i$, $\mathbf{X}_j$, $\mathbf{M}_{ij} = \mathbf{M}_{ji}^\top$, $\mathbf{C}_{10}$ can be expressed in function of $\mathbf{C}_{N0}$ as

$$\mathbf{C}_{10} = \mathbf{D}_0^{-1} \mathbf{C}_{N0}^\top \mathbf{D}_0 . \tag{41}$$

In general, we will have two kinds of messages, traveling forward and backward in the loop, that can be written as a function of the message itself (same link and same direction) $N+1$ time steps ago. For indices $k \in \{0, 1, \ldots, N\}$, with $k \pm 1$ defined modulo $N+1$, we have

$$\mathbf{m}_{k-1,k}^{(t+N+1)} = \alpha_f \mathbf{C}_{k-1,k} \mathbf{m}_{k-1,k}^{(t)} , \tag{42a}$$

$$\mathbf{m}_{k+1,k}^{(t+N+1)} = \alpha_b \mathbf{C}_{k+1,k} \mathbf{m}_{k+1,k}^{(t)} , \tag{42b}$$

with

$$\mathbf{C}_{k-1,k} = \mathbf{M}_{k-1,k}\mathbf{D}_{k-1}\mathbf{M}_{k-2,k-1}\mathbf{D}_{k-2} \cdot \cdots \cdot \mathbf{M}_{k,k+1}\mathbf{D}_k , \tag{43a}$$

$$\mathbf{C}_{k+1,k} = \mathbf{M}_{k+1,k}\mathbf{D}_{k+1}\mathbf{M}_{k+2,k+1}\mathbf{D}_{k+2} \cdot \cdots \cdot \mathbf{M}_{k,k-1}\mathbf{D}_k , \tag{43b}$$

where the similarity transformation

$$\mathbf{C}_{k+1,k} = \mathbf{D}_k^{-1} \mathbf{C}_{k-1,k}^\top \mathbf{D}_k \tag{44}$$

holds.

B. Contraction–based Convergence Analysis

From (32) we have that $\left(\mathbf{b}_k^{(t+n(N+1))}\right)_{n\in\mathbb{N}}$ converges if the sequences $\left(\mathbf{m}_{k-1,k}^{(t+n(N+1))}\right)_{n\in\mathbb{N}}$, $\left(\mathbf{m}_{k+1,k}^{(t+n(N+1))}\right)_{n\in\mathbb{N}}$ in (42a) and (42b) converge. Now recall that the ψ_{ij}'s and ψ_i's are nonnegative valued functions. It follows that the entries of $\mathbf{C}_{k-1,k}$ and $\mathbf{C}_{k+1,k}$ are nonnegative. The following theorem is an immediate consequence of Theorem 8.2 and establishes convergence of belief propagation as a consequence of contraction of the Hilbert metric under the action of a positive linear operator.

Theorem 9.1: Consider the hidden reciprocal model in Figure 3 with the $\mathbf{X}_k$'s taking values in the finite alphabet $\mathcal{X} = \{0, 1, \dots, D-1\}$ and denote by $\mathbf{v}_f$ and $\mathbf{w}_f$ the principal eigenvector of $\mathbf{C}_{k-1,k}$ and $\mathbf{C}_{k+1,k}$, respectively. If the matrices $\mathbf{C}_{k-1,k}$ and $\mathbf{C}_{k+1,k}$ are primitive, then the messages $\mathbf{m}_{k-1,k}$ and $\mathbf{m}_{k+1,k}$ in (42a), (42b) converge in direction to $\mathbf{v}_f$ and $\mathbf{w}_f$, respectively, and the belief at node $\mathbf{X}_k$ converges to $\mathbf{b}_k = \mathbf{m}_{kk} \odot \mathbf{v}_f \odot \mathbf{w}_f$. The rate of convergence is at least linear.

C. Necessary and sufficient conditions for convergence

Many necessary and sufficient conditions for asymptotic stability (e.g. the Jury criterion, the Lyapunov theorem) become simpler in the case of positive systems. From the theory of linear positive systems (see e.g. [29], [2], [16] for details), the following necessary and sufficient conditions for convergence of belief propagation for reciprocal processes follow.

Theorem 9.2: Consider the message update (42a) and denote by $\Lambda(\mathbf{C}_{k-1,k})$ the spectrum of $\mathbf{C}_{k-1,k}$. The following are equivalent:

(i) the iteration (42a) converges ($|\Lambda(\mathbf{C}_{k-1,k})| < 1$);

(ii) all the leading principal minors of the matrix $\mathbf{I} - \mathbf{C}_{k-1,k}$ are positive;

(iii) the coefficients of the characteristic polynomial of $\mathbf{C}_{k-1,k} - \mathbf{I}$ are positive;

(iv) there exists a diagonal matrix $\mathbf{P}$ with positive diagonal elements such that the matrix $\mathbf{C}_{k-1,k}^\top \mathbf{P} \mathbf{C}_{k-1,k} - \mathbf{P}$ is negative definite.

D. Accuracy of the approximation

So far, we have been dealing with *convergence* of the sequence of the beliefs $\left(\mathbf{b}_k^{(t+n(N+1))}\right)_{n\in\mathbb{N}}$. This Section is about *accuracy* of the approximation. The relationship between the posterior probabilities estimated via the belief propagation algorithm on a single loop network and the actual posteriors has been examined in [38], and the analysis is directly applicable to the case

of reciprocal processes. In particular, it has been shown that the smaller is the ratio between the subdominant and the dominant eigenvalue, the smaller is the approximation error. The result is reported here for the sake of completeness.

Theorem 9.3: [38] Consider the reciprocal chain in Figure 3 with the $\mathbf{X}_k$, $k = 0, \ldots, N$ taking values in $\mathcal{X} = \{0, 1, \ldots, D-1\}$. Denote by $\lambda_1, \lambda_2, \ldots, \lambda_D$ the eigenvalues of $\mathbf{C}_{k-1,k}$ (equivalently, see (44), of $\mathbf{C}_{k+1,k}$) sorted by decreasing magnitude and denote by $\mathbf{C}_{k+1,k} = \mathbf{S}\Lambda\mathbf{S}^{-1}$ the eigen-decomposition of $\mathbf{C}_{k+1,k}$. Then the steady-state belief $\mathbf{b}_k$, $k = 0, \ldots, N$ is related to the correct posterior marginal $\mathbf{p}_k$ by:

$$\mathbf{b}_k = \beta \mathbf{p}_k + (1-\beta)\mathbf{q}_k \tag{45}$$

where β is the ratio of the largest eigenvalue of $\mathbf{C}_{k-1,k}$ to the sum of all eigenvalues, $\beta = \lambda_1 / \left(\sum_{j=1}^{D} \lambda_j\right)$, and (the i–th component of the vector) $\mathbf{q}_k$ is given by

$$\mathbf{q}_k(i) = \frac{\sum_{j=2}^{D} \mathbf{S}(i,j)\lambda_j \mathbf{S}^{-1}(j,i)}{\sum_{j=2}^{D} \mathbf{S}(i,j)\lambda_j} .$$

Following [38], we note the fundamental role played by the ratio between the subdominant and the dominant eigenvalue: when this ratio is small, loopy belief propagation converges rapidly and the approximation error is small. Indeed from (45) we have

$$\mathbf{p}_k - \mathbf{b}_k = (1-\beta)\left(\mathbf{q}_k + \mathbf{p}_k\right) ,$$

i.e. the error is small when the maximum eigenvalue dominates the eigenvalue spectrum.

In [38], local correction formulas that compute the correct posteriors on the basis of locally available information have also been provided. In particular, in the case of binary latent variables, it has been shown that

$$p_k(\mathbf{x}_i) = \frac{\lambda_1 b_k(\mathbf{x}_i) + \lambda_2 (1 - b_k(\mathbf{x}_i))}{\lambda_1 + \lambda_2}$$

so that

$$p_k(\mathbf{x}_i) - p_k(\mathbf{x}_j) = \frac{\lambda_1 - \lambda_2}{\lambda_1 + \lambda_2} \left(b_k(\mathbf{x}_i) - b_k(\mathbf{x}_j)\right) \tag{46}$$

i.e. $p_k(\mathbf{x}_i) - p_k(\mathbf{x}_j)$ is positive if and only if $b_k(\mathbf{x}_i) - b_k(\mathbf{x}_j)$ is positive, namely the calculated beliefs are guaranteed to be on the correct side of 0.5, and the correct posterior at node $\mathbf{X}_k$, $\mathbf{p}_k$, can be obtained from the estimated belief $\mathbf{b}_k$ as

$$\mathbf{p}_k = \frac{1}{1+r}\mathbf{b}_k + \frac{r}{1+r}\left(1 - \mathbf{b}_k\right) , \tag{47}$$

where $r := \lambda_2/\lambda_1$ and all eigenvectors and eigenvalues can be calculated by performing operations on the messages that a node receives (see [38] for details). We refer the reader to [38, Section 4.2] for a simulation example where loopy belief propagation has been applied to a single loop network (the hidden reciprocal model in Figure 3).

X. Conclusions

In this paper, we have provided a probabilistic graphical model for reciprocal processes. In particular, it has been shown that a reciprocal process admits a single loop undirected graph as a perfect map. While in the literature a significant amount of attention has been focused on developing dynamical models for reciprocal processes, probabilistic graphical models for reciprocal processes have not been considered before. This approach is distribution independent and leads to a principled solution of the smoothing problem via message passing algorithms for graphical models. For the finite state space case, convergence analysis has been revisited leveraging on contraction properties of positive operators with respect to the Hilbert metric, an argument that is geometric in nature and as such can be extended to study convergence of message passing algorithms to more general settings (state–spaces and graph topologies, see the companion paper [7], where convergence of belief propagation for Gaussian reciprocal processes has been addressed).

References

[1] R. Ackner and T. Kailath. Discrete-time complementary models and smoothing. *International Journal of Control*, 49(5):1665–1682, 1989.

[2] A. Berman and R. J. Plemmons. *Nonnegative matrices in the Mathematical Sciences*. SIAM Press, Philadelphia, 1994.

[3] S Bernstein. Sur les liaisons entre les grandeurs aléatoires. *Verh. Internat. Math.-Kongr., Zurich*, pages 288–309, 1932.

[4] G. Birkhoff. Extensions of Jentzch's Theorem. *Trans. Amer. Math. Soc.*, 85:219–227, 1957.

[5] C. Bishop. *Pattern recognition and machine learning*. Springer, 2006.

[6] P.J. Bushell. Hilbert's metric and positive contraction mappings in a Banach space. *Archive for Rational Mechanics and Analysis*, 52(4):330–338, 1973.

[7] F. P. Carli. On the geometry of message passing algorithms for Gaussian reciprocal processes. *arXiv:1603.09279*, 2016.

[8] F. P. Carli, A. Ferrante, M. Pavon, and G. Picci. A maximum entropy solution of the covariance extension problem for reciprocal processes. *IEEE Trans. on Automatic Control*, 56(9):1999–2012, 2011.

[9] F. P. Carli, A. Ferrante, M. Pavon, and G. Picci. An efficient algorithm for maximum entropy extension of block-circulant covariance matrices. *Linear Algebra and its Applications*, 439(8):2309–2329, 2013.

[10] J-P. Carmichael, J-C. Massé, and R. Theodorescu. Processus gaussiens stationnaires réciproques sur un intervalle. *CR Acad. Sci. Paris Sér. I Math*, 295(3):291–293, 1982.

[11] F. Carravetta. Nearest-neighbor modelling of reciprocal chains. *Stochastics: An International Journal of Probability and Stochastics Processes*, 80(6):525–584, 2008.

[12] F. Carravetta and L. B. White. Modelling and estimation for finite state reciprocal processes. *IEEE Transactions on Automatic Control*, 57(9):2190–2202, 2012.

[13] D. A. Castañon, B.C. Levy, and A.S. Willsky. Algorithms for the incorporation of predictive information in surveillance theory. *International journal of systems science*, 16(3):367–382, 1985.

[14] G. Debreu and I.N. Herstein. Nonnegative square matrices. *Econometrica*, pages 597–607, 1953.

[15] A. P. Dempster. Covariance selection. *Biometrics*, 28:157–175, 1972.

[16] L. Farina and S. Rinaldi. *Positive linear systems: theory and applications*, volume 50. John Wiley & Sons, 2000.

[17] M. Golumbic. *Algorithmic Graph Theory and Perfect Graphs*. Academic Press, New York, 1980.

[18] D. Hilbert. Über die gerade linie als kürzeste verbindung zweier punkte. *Mathematische Annalen*, 46(1):91–96, 1895.

[19] A. T. Ihler, J. Fisher, and A. S. Willsky. Loopy belief propagation: Convergence and effects of message errors. *Journal of Machine Learning Research*, 6:905–936, 2005.

[20] B. Jamison. Reciprocal Processes: The stationary Gaussian case. *The Annals of Mathematical Statistics*, 41:1624–1630, 1970.

[21] B. Jamison. Reciprocal processes. *Probability Theory and Related Fields*, 30(1):65–86, 1974.

[22] B. Jamison. The Markov processes of Schroedinger. *Probability Theory and Related Fields*, 32(4):323–331, 1975.

[23] E. Kohlberg and J.W. Pratt. The contraction mapping approach to the Perron–Frobenius theory: Why Hilbert's metric? *Mathematics of Operations Research*, 7(2):198–210, 1982.

[24] D. Koller and N. Friedman. *Probabilistic graphical models: principles and techniques*. MIT press, 2009.

[25] A. J. Krener, R. Frezza, and B. C. Levy. Gaussian reciprocal processes and self-adjoint stochastic differential equations of second order. *Stochastics and stochastic reports*, 34(1-2):29–56, 1991.

[26] A.J. Krener. Reciprocal diffusions and stochastic differential equations of second order. *Stochastics*, 24(4):393–422, 1988.

[27] S.L. Lauritzen. *Graphical models*. Oxford University Press, 1996.

[28] B.C. Levy, R. Frezza, and A.J. Krener. Modeling and estimation of discrete-time gaussian reciprocal processes. *IEEE Transactions on Automatic Control*, 35(9):1013–1023, 1990.

[29] D. Luenberger. *Introduction to dynamic systems: theory, models, and applications*. Wiley, 1979.

[30] C. C. Moallemi and B. Van Roy. Convergence of min-sum message passing for quadratic optimization. *IEEE Transactions on Information Theory*, 55(5):2413–2423, 2009.

[31] J. M. Mooij and H. J. Kappen. Sufficient conditions for convergence of the sum–product algorithm. *IEEE Transactions on Information Theory*, 53(12):4422–4437, 2007.

[32] J. Pearl. *Probabilistic reasoning in intelligent systems: Networks of plausible reasoning*. Morgan Kaufmann Publishers, 1988.

[33] J. Pearl and A. Paz. Graphoids: A graph-based logic for reasoning about relevance relations. Technical report, University of California (Los Angeles), 1985.

[34] G. Picci and F. P. Carli. Modelling and simulation of images by reciprocal processes. In *Proc. of the Tenth International Conference on Computer Modeling and Simulation, UKSIM 2008*, pages 513–518, 2008.

[35] J. A. Sand. Reciprocal realizations on the circle. *SIAM J. Control and Optimization*, 34:507–520, 1996.

[36] L. Srinivasan, U.T. Eden, A.S. Willsky, and E.N. Brown. A state-space analysis for reconstruction of goal-directed movements using neural signals. *Neural computation*, 18(10):2465–2494, 2006.

[37] M. J. Wainwright and M. I. Jordan. Graphical models, exponential families, and variational inference. *Foundations and Trends® in Machine Learning*, 1(1-2):1–305, 2008.

[38] Y. Weiss. Correctness of local probability propagation in graphical models with loops. *Neural computation*, 12(1):1–41, 2000.

[39] Y. Weiss and W.T. Freeman. Correctness of belief propagation in gaussian graphical models of arbitrary topology. *Neural computation*, 13(10):2173–2200, 2001.

[40] L.B. White and F. Carravetta. Optimal smoothing for finite state hidden reciprocal processes. *IEEE Transactions on Automatic Control*, 56(9):2156–2161, 2011.